



Scientific Linux

The Experiment That Worked

Troy Dawson
dawson@fnal.gov

LinuxWorld Russia 2006
September 5, 2006

Scientific Linux

- Presentation Overview
 - What is Scientific Linux
 - Why was S.L. 'Experimental'
 - What we did to help the experiment
 - How did it succeed
 - Conclusion



What Is Scientific Linux?

- Recompiled Red Hat Enterprise Linux
- Value Added
 - Changes
 - Additions
 - packages
 - tweaks

What Is Scientific Linux?

- Changes
 - Where we legally must because of Red Hat's trademarks
 - Where we morally felt we should
 - Bookmarks, Up2date (automated patch updates)
 - Other changes discussed later
- Additions
 - Additions discussed later



Why Was S.L. Experimental?

- At the time, all RHEL recompiled distributions were experimental
- For the High Energy Physics community, there had never been a single software package that everyone agreed upon.

Helping the Experiment

How Did We Help S.L. Succeed

- Additions and Changes
- Ability to “sit” on a release
- A community infrastructure



Helping the Experiment

Additions and Changes

- Additions
 - Packages
 - Tweaks
- Changes
 - Packages
 - Installer
 - Sites

Helping the Experiment

Additions and Changes

- Additions
 - Packages
 - Consolidated RHEL suites
 - GFS, Cluster Suite
 - Commonly used in science
 - OpenAFS
 - That we felt were missing
 - pine, perl-CPAN, icewm



Helping the Experiment

Additions and Changes

- Additions

- Tweaks

- Easily installed
 - Don't have to modify original package
 - Not installed by default
 - Modifies standard configuration's
 - color ls
 - Turns off color ls
 - terminal button on desktop panel
 - security enhancements
 - serial console output



Helping the Experiment

Additions and Changes

- Changes
 - Where we legally must because of Red Hat's trademark's
 - Where we morally felt we should.
 - Bookmarks, Up2date (automated patch updates)
 - Installation program modified for Sites



Helping the Experiment Sites

- What are Sites?
- Why did we create the ability to do sites?
- How do you build a site?

Helping the Experiment

What are sites

- Sites are a way for a Laboratory, or site, to customize Scientific Linux with minimum effort and changes
- Gives facilities choices
 - What packages are installed
 - Add, delete, modify packages
 - Customize installer program




Helping the Experiment

What are sites

To Summarize

Sites allow you to change Scientific Linux into your own custom distribution.



Helping the Experiment

Why Sites

- Each Lab had its own environment
 - Time Zone
 - Authentication Policies
 - Kerberos 5, Kerberos 4, LDAP
 - Extra or Different Packages
 - Customized Kernels, Better Multimedia
 - Their own update servers



Helping the Experiment

Why Sites

- Each Lab already had to customize whichever Linux distribution they used
- We hoped that with a distribution that is easy to customize, the labs would unify on one distribution - Scientific Linux

Helping the Experiment

How to build a site

- In each release there is a sites directory for each architecture
 - Example: `scientific/43/i386/sites/`
- There is an example site for each area
 - Example: `scientific/43/i386/sites/example/`
- There is documentation on how to build a site, in each example site
 - `scientific/43/i386/sites/example/build/documentation/`



Helping the Experiment

How to build a site

- For our example, we'll call our new site 'newsite'
- First: copy the entire release to a local area
 - `rsync -avH --exclude=errata/`
`rsync://rsync.scientificlinux.org/scientific/43/i386/ /scientific/43/i386/`
 - You can exclude the whole errata directory, for more information
<https://www.scientificlinux.org/download/mirroring/mirror.rsync>
- Second: copy the example site to newsite
 - `cp -r /scientific/43/i386/sites/example/`
`/scientific/43/i386/sites/newsite/`



Helping the Experiment

How to build a site

- Add any packages you want
 - `cd /scientific/43/i386/sites/newsite/Updates`
 - Add whatever rpm's you wish here
- Change what packages get installed
 - `cd /scientific/43/i386/sites/newsite/base`
 - Edit `comps.xml.main`
 - Edit `scripts/comps.site.spec`
 - `sh scripts/cp.comps.xml.main.sh`
- Change the installation pictures
 - `cd /scientific/43/i386/sites/newsite/RHupdates/pixmaps`
 - Add modified graphics (example: `splash.png`)



Helping the Experiment

How to build a site

- Note that the steps outlined here must be executed on a system which is running Scientific Linux. The system must have the anaconda and anaconda-runtime rpms installed.
- Build our new site
 - `cd /scientific/43/i386/sites/newsite/build/scripts`
 - Edit `locations.include`
 - `TREE=/scientific/43/i386`
 - `SITE=newsite`
 - `cp rpmdb-example.spec.in rpmdb-newsite.spec.in`
 - Then edit it and change `example` to `newsite`
 - `sh make.everything.sh`



Helping the Experiment

How to build a site

- Network Install images
 - `cd /scientific/43/i386/sites/newsite/images/`
 - When doing a network install, you would point the installer to `/scientific/43/i386`
- CD Install ISO images
 - `cd /scientific/43/i386/sites/newsite/iso/`

Helping the Experiment

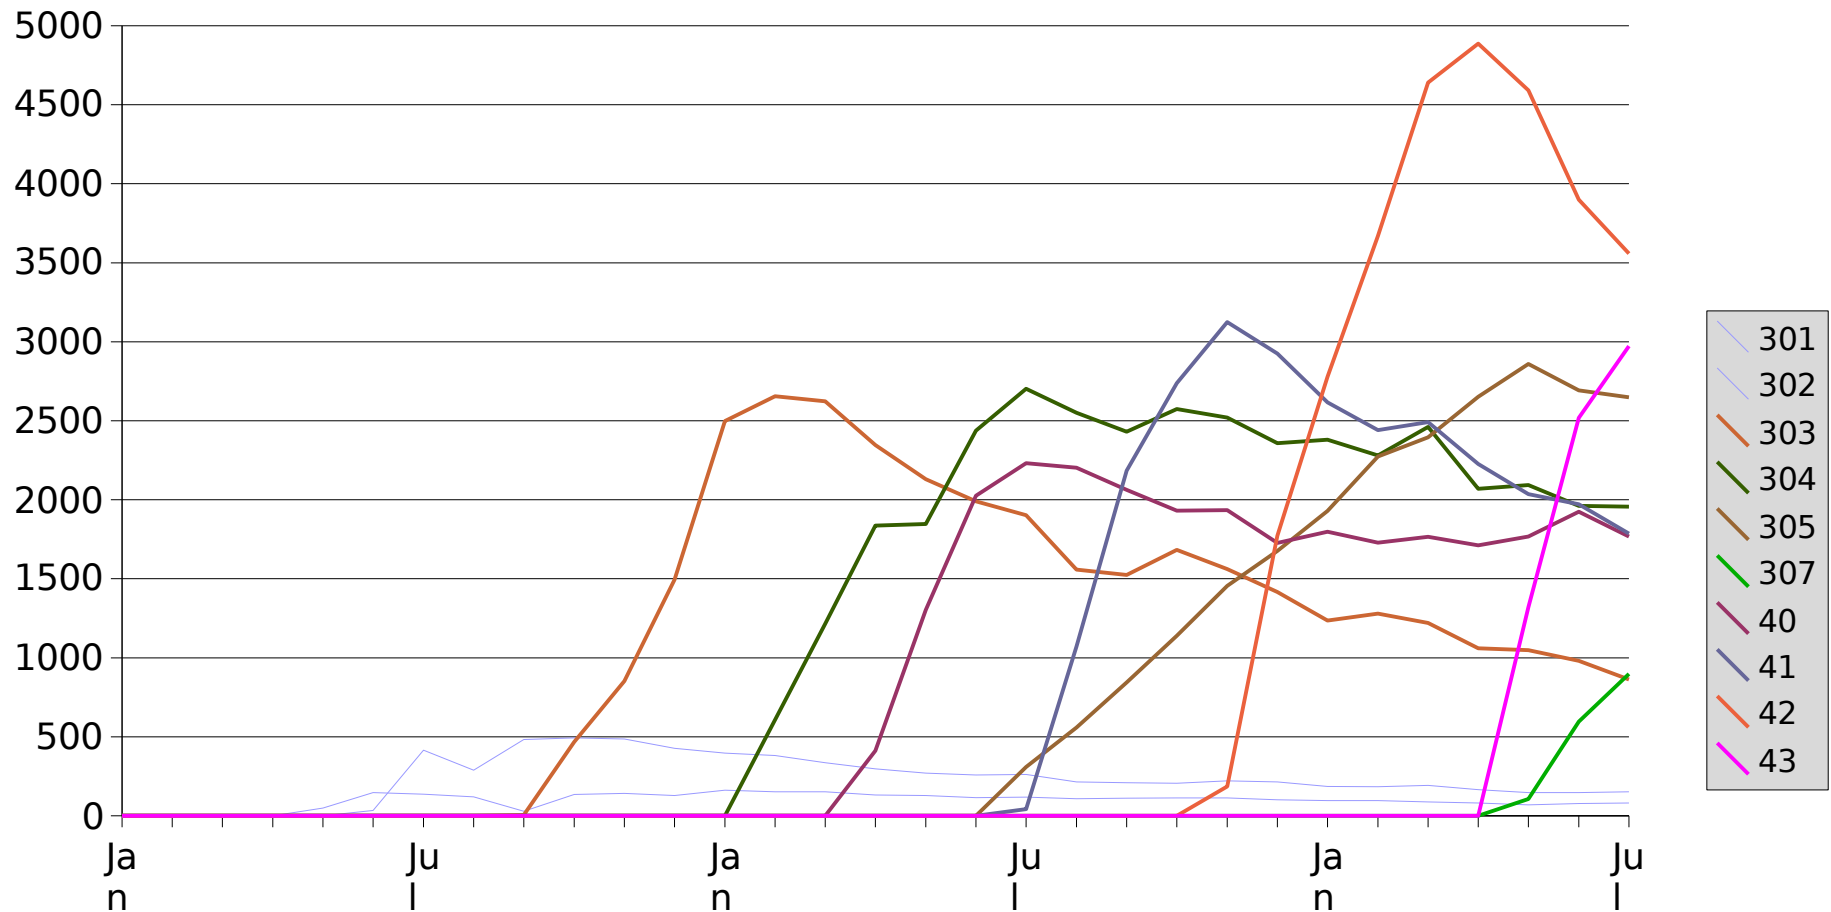
Ability To Sit On A Release

- One of the goals for Scientific Linux was that scientists would be able to install one version of Scientific Linux, and not have to worry about upgrading.
 - They would still get security updates
 - For a multi-year experiment, this is very important. Changes in libraries might change the result a an experiment



Helping the Experiment Ability To Sit On A Release

2004 to Present - By Individual Release



Helping the Experiment

Ability To Sit On A Release

- Red Hat currently does not let a user do this. All security updates state that the user must be running the latest updates.
- From what I know, no other Recompiled RHEL distribution does this either.
- The effort and storage is high



Helping the Experiment

Ability To Sit On A Release

<https://www.scientificlinux.org/download/mirroring/mirror.size>

S.L. 3.0.x

301 - 9G

302 - 27G

303 - 23G

304 - 26G

305 - 28G

307 - 11G

S.L. 4.x

41 - 23G

42 - 17G

43 - 12G



Scientific Linux

Helping the Experiment

A Community Infrastructure

- Mail Lists
- Web Site
- File Servers
- Semi-Annual Meetings

Helping the Experiment

A Community Infrastructure

- Mail Lists
 - <https://www.scientificlinux.org/maillists/>
 - scientific-linux-users@fnal.gov
 - scientific-linux-devel@fnal.gov
 - scientific-linux-errata@fnal.gov
 - scientific-linux-announce@fnal.gov

Helping the Experiment

A Community Infrastructure

- Web Site
 - <https://www.scientificlinux.org/>
- File Servers
 - <https://www.scientificlinux.org/download/>
 - <ftp://ftp.scientificlinux.org/linux/scientific/>
 - <rsync://rsync.scientificlinux.org/scientific/>
 - 22 public mirrors



Helping the Experiment

A Community Infrastructure

- Semi-Annual Meetings
 - <https://www.hepix.org/>
 - High Energy Physics (HEP) Unix users and administrators collaboration
 - For the past two years, the last morning of the conference has been for steering the direction of Scientific Linux

Scientific Linux

Success

- Adoption by the HEP scientific community
- Adoption by non-HEP users
- Collaboration between Labs
 - Adding more lead developers
- Community answering it's own problems
- Looking at the numbers



Scientific Linux

Success

- Adoption by the HEP scientific community
 - This happened at our first release S.L. 3.0.1
 - Very unexpected
 - One lab even switched from SUSE to S.L.

Scientific Linux

Success

- Adoption by non-HEP users
 - We never actively advertised
 - Universities associated with physics started using it
 - Several people outside the scientific and educational community wrote about us
 - I am constantly running into places where Scientific Linux is being used.



Scientific Linux

Success

- Collaboration between Labs
 - With the initial release, S.L. 3.0.1, CERN officially started collaborating with us.
 - Later, DESY took over the lead of openAFS.
 - This year, PSI started doing the Scientific Linux Live CD and DVD.
 - Each developer that takes on a project, helps reduce the workload.



Scientific Linux

Success

- Community answering it's own problems
 - When this happened, that was when I said we had succeeded.
 - Initially, I made sure questions to scientific-linux-users was answered fast and politely. Soon, more people started joining the list, and found it was a good place for information
 - When a critical mass was reached, I found that I no longer had to answer the questions, and that often, there were better answers.



Scientific Linux


Success

- The following statistics were gathered from <ftp.scientificlinux.org> log files
 - These numbers are the minimum. We know that they are higher.
 - These numbers are based on the yum and apt access to the errata

Scientific Linux

Success

2004 to Present - By Release




Scientific Linux

Scientific Linux

Success


2004 to Present - By Individual Release



Scientific Linux

Scientific Linux

Top 10 Downloads by Country



Scientific Linux

Scientific Linux

Conclusion

- Scientific Linux started out with the question “If we build this, will anyone use it.”
- We built it with a lot of care and some unique features that administrators of large sites would want.
- Scientific Linux has been very successful, in every way that we look at it.



Scientific Linux

References

- <https://www.scientificlinux.org/>
- <ftp://ftp.scientificlinux.org/linux/scientific/>